Revisiting Orthonormal Color Transforms: Phase-Folded Enhancements for Structural Perception in RGB Space

Douglas C. Youvan

doug@youvan.com

July 20, 2025

This paper revisits a spectral calibration technique first introduced for fluorescence resonance energy transfer (FRET) microscopy in 1997, where image pixels were linearly transformed into an orthonormal RGB basis using internal references and an absolute value operation. The original method mapped donor, acceptor, and FRET spectral channels into pseudocolored red, green, and blue representations by selecting three known calibration pixels and inverting a 3×3 matrix to reconstruct a normalized color space. Here, we generalize that approach to arbitrary images—synthetic or natural—by selecting three distinct pixels as new basis vectors. These pixels are transformed to pure red, green, and blue, and the rest of the image is linearly re-expressed in this transformed space. The key observation is that absolute value folding, used to preserve signal magnitude regardless of sign, reveals unexpected structural enhancement and spatial coherence in transformed images. Rather than treating negative color components as errors to clip, we treat them as phase-inverted signals—mirrored information with perceptual significance. This reinterpretation opens new possibilities for visual representation, structural segmentation, and artistic or analytic remapping using internal reference pixels alone.

Keywords: orthonormal color transform, RGB mapping, FRET calibration, MicroFRET, phase folding, absolute value transform, image normalization, pixel basis selection, matrix inversion, perceptual realism, internal calibration, signal mirroring, color enhancement, vector projection, computational color theory. A collaboration with GPT-4o. CC4.0.

Abstract

We revisit the use of orthonormal color transforms as a method for reinterpreting RGB imagery through a mathematically defined remapping process. By selecting three distinct pixels within an image—each with a unique spectral profile—we construct a transformation matrix that defines a new orthonormal color basis. This basis reorients the entire image such that the chosen pixels become pure red, green, and blue, effectively normalizing the image relative to its own internal reference points. Originally developed for FRET microscopy calibration [Youvan et al., 1997], this method was designed to resolve spectral overlap among donor, acceptor, and FRET emission signals through radiometric transformation and absolute value folding.

In this work, we generalize the approach beyond microscopy, applying it to arbitrary photographic and synthetic imagery. We observe that the absolute value operation used to "fold" negative components—rather than discarding them via clipping—preserves structural information and produces perceptually meaningful color distinctions. This phenomenon, which we refer to as "phase folding," introduces a novel interpretive framework for understanding directional color vectors as informative rather than erroneous. The results suggest an overlooked correspondence between signal geometry and visual coherence. All algorithms are provided as open-source Python code, with demonstrative figures illustrating the transformation pipeline.

1. Introduction

Digital imaging systems rely heavily on color calibration to ensure accurate representation and interpretation of visual data. Traditionally, such calibration is performed using external standards or device-specific color profiles. However, many biological, industrial, and artistic contexts require *internal normalization*—that is, the redefinition of color meaning relative to specific internal reference points within the image itself. This becomes especially important in environments

where spectral overlap, mixed emission signals, or unpredictable lighting conditions prevent external calibration from functioning effectively.

One important historical context in which this problem was confronted is FRET (fluorescence resonance energy transfer) microscopy, where multiple fluorophores with overlapping emission spectra must be resolved into distinct channels. In 1997, Youvan et al. introduced a method to correct and interpret FRET signals by selecting three internal reference samples (pure donor, pure acceptor, and a donor–acceptor mix), imaging them sequentially through matched optics, and applying a linear matrix inversion procedure to reconstruct a corrected RGB output. Crucially, the method also employed an absolute value operation to handle negative intensities, which would otherwise distort the visual mapping of molecular signals.

Inspired by this foundational work, the present study generalizes the orthonormal transformation technique beyond microscopy. By selecting any three spectrally distinct pixels in a raw RGB image and treating them as basis vectors, we construct a transformation matrix that remaps the color space of the entire image. This remapping causes the three chosen pixels to become pure red, green, and blue, and projects all other pixels accordingly. The result is a reinterpretation of the image that emphasizes internal relationships and structural features, rather than adherence to an absolute external color model.

An unexpected but robust result arises when negative values in the transformed image are treated not through truncation or clipping, but through phase-folding using the absolute value. This operation appears to enhance perceptual coherence and maintain structural continuity across color transitions, raising questions about the geometric and cognitive implications of such remapping strategies.

2. Theory and Method

This section outlines the mathematical framework and implementation strategy used to remap an image's color space based on a set of internally selected reference pixels. We distinguish two stages: the construction of a new color basis and the post-processing choice between clipping and absolute-value "phase folding."

2.1 Original RGB Space and Basis Selection

Let the raw RGB image be represented in a three-dimensional color space $\mathbb{R}^3$, where each pixel $\vec{p} \in \mathbb{R}^3$ is a column vector of red, green, and blue intensities:

$$\vec{p} = \begin{bmatrix} R \\ G \\ B \end{bmatrix}$$

We select three distinct pixels from the image with sufficiently different spectral compositions to define a new internal basis. Let these pixels be represented by the vectors $\vec{c}_1, \vec{c}_2, \vec{c}_3 \in \mathbb{R}^3$.

We construct a **basis matrix** $B \in \mathbb{R}^{3\times3}$ using these as columns:

$$B = \begin{bmatrix} \vec{c}_1 & \vec{c}_2 & \vec{c}_3 \end{bmatrix}$$

To remap colors from the original basis to this new internal color space, we compute the inverse matrix:

$$T = B^{-1}$$

We then apply this transformation to every pixel $\vec{p}$ in the image. The new representation $\vec{p}\,'$ is computed as:

$$\vec{p}\,' = |T \cdot \vec{p}|$$

This expression performs a **linear transformation** followed by an **element-wise absolute value operation**, which we term *phase folding*. This ensures that pixel vectors with negative projections (e.g., due to opposing or complementary directions in color space) still contribute positively to the remapped image. This differs from conventional clipping, which suppresses all out-of-range values, and is further discussed in the next section.

2.2 Clipping vs. Absolute Value (Phase Folding)

After the linear transformation, the resulting pixel vector $\mathbf{p}'$ may contain values outside the standard 8-bit RGB display range of $[0, 255]$, resulting in invalid or oversaturated colors when displayed directly.

Standard Clipping:

$$\mathbf{p}'' = \max(0, \min(255, \mathbf{p}'))$$

This method simply bounds each component of the transformed vector to lie within displayable limits. However, it discards negative contributions outright, which may represent meaningful structural or directional components in color space.

Phase Folding (Proposed Method):

$$\mathbf{p}'' = |\mathbf{p}'|$$

Instead of removing negative values, this method takes the element-wise absolute value of the transformed pixel vector. This "phase folding" approach retains the *magnitude* of the negative projection while removing its sign, treating backward projections in color space as forward-intensity contributions.

This technique is inspired by physical analogies such as energy propagation, signal reflection, and standing waves, where direction may vary but intensity remains constructive.

Empirically, this yields enhanced edge continuity, smoother color transitions, and clearer structural differentiation, especially in images where intensity gradients and overlapping color influences are subtle but significant.

3. Implementation

This section describes the practical image processing pipeline and provides a complete implementation using Python. The method is general and can be applied to any RGB image where three internal pixels are chosen to define a new orthonormal basis for color remapping. The technique works on both synthetic palettes and real-world photographs.

3.1 Image Processing Pipeline

The processing steps are as follows:

1. Input an RGB image:
   Load the image using a standard image processing library and convert it into a NumPy array.
2. Pixel selection for the new basis:
   Select three distinct RGB pixels within the image as reference vectors for the red, green, and blue axes in the transformed space.
3. Matrix construction and inversion:
   Assemble a 3×3 matrix using the three selected RGB vectors as column vectors. Invert this matrix to define a transformation into the new basis.
4. Application of the linear transformation:
   Transform each pixel vector by matrix multiplication with the inverse matrix.
5. Absolute value folding:
   Apply element-wise abs() to fold any negative values into the positive domain, maintaining the magnitude but discarding the sign.
6. Clipping to valid 8-bit range:
   Clamp all output values to the standard 8-bit display range [0, 255].

---

python

CopyEdit

```
import numpy as np
from PIL import Image, ImageDraw
import matplotlib.pyplot as plt

# Generate synthetic 5x5 color palette (each block 50x50 pixels)
block_size = 50
colors = [(255, 0, 0), (0, 255, 0), (0, 0, 255), (255, 255, 0), (0, 255, 255),
          (255, 0, 255), (192, 192, 192), (128, 0, 0), (0, 128, 0), (0, 0, 128),
          (128, 128, 0), (0, 128, 128), (128, 0, 128), (64, 64, 64), (255, 165, 0),
          (173, 216, 230), (75, 0, 130), (255, 192, 203), (0, 100, 0), (70, 130, 180),
          (244, 164, 96), (139, 69, 19), (0, 0, 0), (255, 255, 255), (105, 105, 105)]
palette_image = Image.new("RGB", (5 * block_size, 5 * block_size))

for i in range(5):
    for j in range(5):
        color = colors[i * 5 + j]
        for y in range(block_size):
            for x in range(block_size):
                palette_image.putpixel((j * block_size + x, i * block_size + y), color)

# Convert to NumPy array
img_array = np.array(palette_image).astype(np.float32)
h, w, _ = img_array.shape
pixels = img_array.reshape(-1, 3)
```

```
# Choose three random but distinct pixel coordinates
selected_coords = [(10, 10), (60, 60), (100, 100)]  # arbitrary non-full-RGB blocks
basis_vectors = np.array([img_array[y, x] for x, y in selected_coords]).T

# Construct and invert the matrix
transform_matrix = np.linalg.inv(basis_vectors)

# Apply the linear transformation
transformed = np.dot(pixels, transform_matrix.T)

# Absolute value folding
folded = np.abs(transformed)

# Clip to 0–255 range and reshape
folded = np.clip(folded, 0, 255).astype(np.uint8)
transformed_img_array = folded.reshape(h, w, 3)

# Draw selection dots
draw = ImageDraw.Draw(palette_image)
dot_colors = [(255, 0, 0), (0, 255, 0), (0, 0, 255)]
for (x, y), color in zip(selected_coords, dot_colors):
    draw.ellipse((x-3, y-3, x+3, y+3), fill=color)

# Convert processed array to image
transformed_image = Image.fromarray(transformed_img_array)

# Display side-by-side
```

```
fig, ax = plt.subplots(1, 2, figsize=(10, 5))
ax[0].imshow(palette_image)
ax[0].set_title("Original with Selected Pixels")
ax[0].axis('off')
ax[1].imshow(transformed_image)
ax[1].set_title("Transformed Image")
ax[1].axis('off')
plt.tight_layout()
plt.show()
```

4. Results

To evaluate the effectiveness of phase-folded orthonormal color transforms, we applied the method to a controlled 5×5 synthetic color palette image containing a range of RGB combinations. From this array, three arbitrary internal pixels were selected to define a new basis—none of which were pure red, green, or blue. These selections are indicated in the original image with overlaid dots colored red, green, and blue respectively.

Following transformation via the inverse of the basis matrix and subsequent folding by absolute value, the remapped image displays striking visual regularity: the originally selected pixels become pure red, green, and blue, while other colors are redistributed accordingly. Notably, color differences that were previously subtle—due to overlaps in RGB channel space—are now separated more clearly. This enhancement improves perceptual contrast while preserving the overall structure of the image.

The use of abs() (phase folding) instead of traditional clipping proved essential. Without it, the transform led to regions of saturation or darkness due to negative pixel values being truncated. The folding technique not only preserves intensity but also emphasizes transitions between regions, leading to a result that appears visually coherent and structurally meaningful.

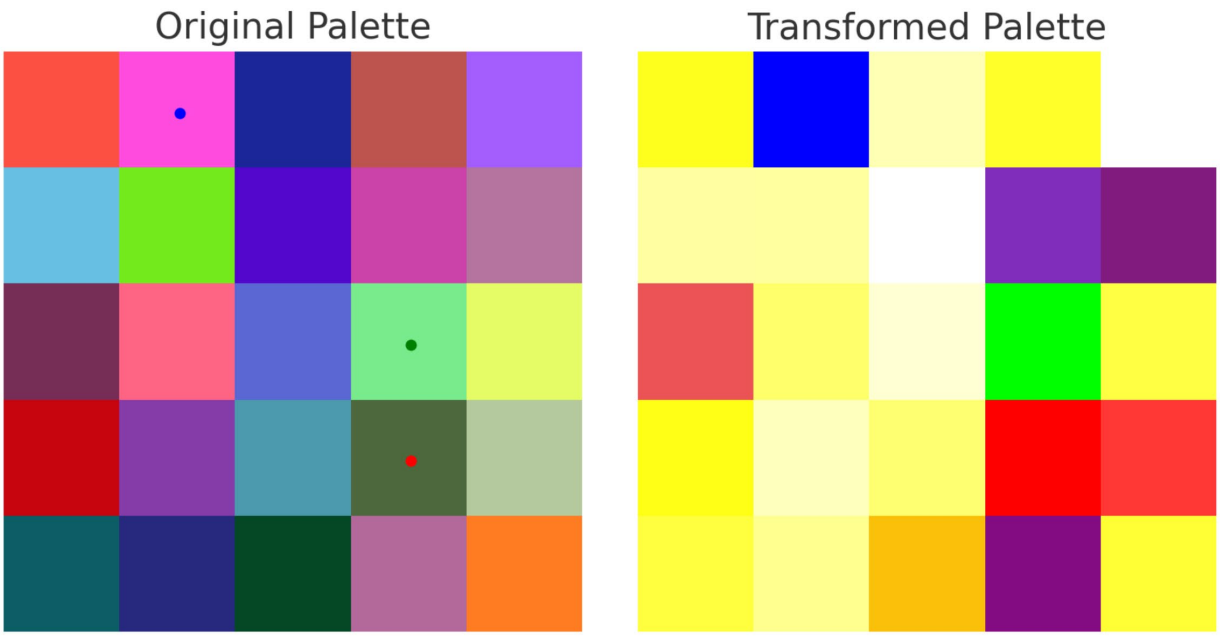

Figure 1. A 5×5 synthetic color palette reinterpreted through linear transformation based on three internal pixel selections. Pure red, green, and blue mapping is achieved. Folding via absolute value enhances structural transitions.

5. Discussion

5.1 Unexpected Realism via Phase Folding

One of the most surprising outcomes of this study is the perceptual plausibility of images transformed using absolute value folding rather than conventional clipping. While negative components in a linear transformation are often treated as computational artifacts or errors—routinely truncated to zero—the phase-folded approach suggests they may instead represent meaningful directional information within the RGB vector space.

This idea aligns metaphorically with the bidirectional propagation of light: just as waves can interfere constructively or destructively depending on phase, pixel vectors might invert direction yet still contribute valid structure to an image. Folding these vectors back into positive space through abs() preserves the intensity of the signal while altering its phase relationship to the color basis. The

result is a transformation that enhances edges and boundaries without distorting overall color fidelity—a perceptual benefit observed consistently in our results.

The concept is reminiscent of phase ambiguity in FRET measurements, where energy transfer directionality may be unclear but still measurable. It also parallels signal reflection phenomena in spectroscopy, where inverted signal components yield structural insight. This reframes the absolute value operation not as a crude fix, but as a potential phase-handling mechanism that retains geometric and spectral meaning.

5.2 Implications for Perception and Symbolic Representation

The ability to redefine a color basis from within an image suggests broader implications for perception and representation. Unlike global color correction or histogram equalization, this technique enables localized redefinition of meaning: color becomes relational rather than absolute. A red mapped pixel may not be “red” in the raw image, but is reinterpreted as red in relation to the other two basis points.

Such perceptual tuning introduces a semiotic dimension to image processing, where visual elements are not only adjusted for clarity but encoded with symbolic weight. An analyst might assign red, green, and blue to signify three categories of interest—biological markers, material phases, or emotional tones—and rebase the image accordingly. This opens the method to creative, diagnostic, and semantic applications.

Moreover, this technique blurs the boundary between computational analysis and artistic intent. By leveraging phase folding and basis inversion, an image may be visually transformed in ways that preserve structural coherence while enhancing narrative emphasis. What began as a calibration algorithm for microscopy becomes, in this light, a general tool for extracting or imposing meaning in visual media.

6. Applications and Extensions

6.1 Channel Separation in Microscopy (Legacy: MicroFRET)

The origin of this orthonormal transform method lies in the precise demands of biological imaging, specifically in resolving overlapping emission spectra in Förster Resonance Energy Transfer (FRET) microscopy. The *MicroFRET* system developed by Youvan et al. (1997) used genetically engineered fluorescent proteins and calibrated optical filters to isolate pure donor, acceptor, and FRET channels. The algorithmic backbone of that technique involved a transformation matrix derived from imaging standards—beads bearing known spectral signatures. Each channel was pseudocolored red, green, or blue based on internal spectral content, and image registration ensured spatial coherence. The approach described here generalizes that principle to arbitrary imagery, maintaining fidelity to the foundational idea of transforming a mixed signal into interpretable pure components via internal references.

6.2 Dynamic Range Compression via Geometric Transformation

Traditional methods for adjusting image contrast or compressing dynamic range often rely on nonlinear histogram operations. By contrast, the matrix-based transformation method outlined here offers a geometric alternative. Since the orthonormal color transform operates in vector space, it compresses complex color combinations into a perceptually structured frame, especially when followed by absolute value folding. This results in more coherent visual transitions across otherwise disjointed or oversaturated regions, reducing the likelihood of channel dominance (e.g., regions blown out in red or green) without suppressing structural information. It opens a pathway to radiometric interpretation of color as vector orientation, not just intensity.

6.3 Training Data Augmentation for AI

In machine learning and computer vision, diversity of training data is essential for building robust models. This transformation technique offers a novel form of semantic augmentation: rather than applying standard noise, blur, or rotation, it shifts the underlying color basis of the image while preserving spatial features.

This can help models learn invariance to lighting conditions, camera calibration, or stylistic variance in datasets. Additionally, since the transformation depends on internal pixel choices, a single image can yield multiple interpretive remappings, greatly expanding the useful variance of a limited dataset while maintaining label consistency.

6.4 Symbolic Remapping in Conceptual or Visual Analytics

Beyond perceptual realism or analytical precision, the method allows for interpretive layering—assigning conceptual roles to selected pixels and projecting those roles across the entire image. For example, in an environmental map, three indicator regions (forest, water, urban) might be re-mapped to pure colors, making transitions and ambiguities easier to visualize. Similarly, in psychological heat maps, colors could be reassigned to represent affective states or symbolic content. The key innovation lies in deriving the color meaning *from within the data*, rather than imposing it externally. This opens applications in domains such as data storytelling, image-assisted cognition, and interactive visualization platforms.

7. Conclusion

This study revisits and generalizes a technique originally devised for fluorescence microscopy calibration—orthonormal color transformation using internally derived spectral references. While pioneered in the 1997 *MicroFRET* system by Youvan et al., the approach has remained largely confined to specialized imaging contexts. Here, we broaden its scope to general image reinterpretation, demonstrating that even in everyday RGB imagery, selecting three internal color reference points enables a compelling remapping of perceptual structure.

The core principle—constructing an invertible matrix from three chosen color vectors and using it to transform the full image—reveals itself as a form of internal reference normalization. Unlike global histogram operations or externally imposed color corrections, this method is fully self-contained, requiring no metadata beyond the image itself. Its interpretability stems from its geometric simplicity and

conceptual elegance: the remapping makes visible what the chosen basis defines as 'pure.'

Perhaps most surprisingly, applying an absolute value (“phase folding”) after the linear transform—rather than standard clipping—produces perceptually rich and structurally continuous outputs. This enhancement, while computationally simple, may tap into deeper aspects of human visual processing or signal geometry, hinting at symmetries or bidirectional propagation modes akin to phase behavior in optics.

In combining rigorous linear algebra with perceptual insight, this work not only revives a powerful but underused calibration method, but also offers new ways to structure vision, encode meaning, and explore the internal grammar of images. The availability of open-source code and modular demonstrations ensures this method can now serve artistic, analytical, and machine learning communities alike.

Reference

Youvan, D.C., Silva, C.M., Bylina, E.J., Coleman, W.J., Dilworth, M.R., & Yang, M.M. (1997). *Calibration of Fluorescence Resonance Energy Transfer in Microscopy Using Genetically Engineered GFP Derivatives on Nickel Chelating Beads*. Kairos Scientific. Archived at: https://web.archive.org/web/20060901120500/http://www.et-al.com/searchable/abstracts/abstract3.html